arXiv:1510.08562v1 [math.OC] 29 Oct 2015

CONVERGENCE RATE OF INCREMENTAL GRADIENT AND INCREMENTAL NEWTON METHODS

M. GÜRBÜZBALABAN*, A. OZDAGLAR* , AND P. PARRILO*

Abstract. The incremental gradient method is a prominent algorithm for minimizing a finite sum of smooth convex functions, used in many contexts including large-scale data processing applications and distributed optimization over networks. It is a first-order method that processes the functions one at a time based on their gradient information. The incremental Newton method, on the other hand, is a second-order variant which exploits additionally the curvature information of the underlying functions and can therefore be faster. In this paper, we focus on the case when the objective function is strongly convex and present fast convergence results for the incremental gradient and incremental Newton methods under the constant and diminishing stepsizes. For a decaying stepsize rule $\alpha_k = \Theta(1/k^s)$ with $s \in (0,1]$, we show that the distance of the IG iterates to the optimal solution converges at rate $\mathcal{O}(1/k^s)$ (which translates into $\mathcal{O}(1/k^{2s})$ rate in the suboptimality of the objective value). For $s > 1/2$, this improves the previous $\mathcal{O}(1/\sqrt{k})$ results in distances obtained for the case when functions are non-smooth. We show that to achieve the fastest $\mathcal{O}(1/k)$ rate, incremental gradient needs a stepsize that requires tuning to the strong convexity parameter whereas the incremental Newton method does not. The results are based on viewing the incremental gradient method as a gradient descent method with gradient errors, devising efficient upper bounds for the gradient error to derive inequalities that relate distances of the consecutive iterates to the optimal solution and finally applying Chung's lemmas from the stochastic approximation literature to these inequalities to determine their asymptotic behavior. In addition, we construct examples to show tightness of our rate results.

1. Introduction.

We consider the following additive cost optimization problem

$$\min \sum_{i=1}^{m} f_i(x) \quad \text{s.t.} \quad x \in \mathbb{R}^n \tag{1.1}$$

where the objective function is the sum of a large number of convex *component functions* $f_i : \mathbb{R}^n \to \mathbb{R}$. Such problems arise in a number of settings including distributed optimization across m agents, where the component function f_i corresponds to the local objective function of agent i [9,23,24,29], and statistical estimation problems where each f_i represents the loss function associated with one of the data blocks [5,8,30,32]. Our goal is to exploit the additive structure of problem (1.1) and solve it using incremental methods which involve sequential processing of component functions.

*Laboratory for Information and Decision Systems, Massachusetts Institute of Technology, Cambridge, MA, 02139, USA. email: {mertg, asuman,parrilo}@mit.edu.

We first consider the *incremental gradient (IG) method* for solving problem (1.1). IG method is similar to the standard gradient method with the key difference that at each iteration, the decision vector is updated incrementally by taking sequential steps along the gradient of the component functions f_i in a cyclic order. Hence, we can view each outer iteration k as a cycle of m *inner iterations*: starting from initial point $x_1^1 \in \mathbb{R}^n$, for each $k \geq 1$, we update the iterate x_i^k as

$$x_{i+1}^k := x_i^k - \alpha_k \nabla f_i(x_i^k), \qquad i = 1, 2, \ldots, m, \tag{1.2}$$

where $\alpha_k > 0$ is a stepsize. We set $x_1^{k+1} = x_{m+1}^k$ and refer to $\{x_1^k\}$ as the *outer iterates*. When the component functions are not smooth, we can replace gradients with a subgradient and the corresponding method is called the incremental subgradient method. Using the update relation (1.2), for each $k \geq 1$, we can write down the relation between the outer iterates as

$$x_1^{k+1} = x_1^k - \alpha_k \sum_{i=1}^{m} \nabla f_i(x_i^k), \tag{1.3}$$

where $\sum_{i=1}^m \nabla f_i(x_i^k)$ is the aggregated component gradients and serve as an approximation to the full gradient $\nabla f(x_1^k)$ with the difference that it is evaluated at different inner iterates.

IG is a prominent algorithm with a long history that has appeared in many contexts. In the artificial intelligence literature, it has been used in training neural networks in the 80s and is known as the online backpropagation algorithm [5, 17, 34]. Another well-known example of this method is the Kaczmarz method for solving systems of linear equations, which is a special case of the IG method [6].

Due to the simplicity and long history of IG method, its global convergence has been supported under various conditions (see [5] for a survey), however characterizing its convergence rate has been the subject of more recent work . Among the papers relevant to our work, Kohonen [16] focused on quadratic component functions with constant stepsize, $\alpha_k = \alpha > 0$ for all k, and showed that the iterates may converge to a limit cycle (subsequence of inner iterates converge to different limits close to optimal). The papers [2, 3, 7, 11, 12, 17–19, 33] focused on diminishing stepsize and showed convergence of the algorithm and its variants under different assumptions. The papers [31] and [22] studied IG with a constant stepsize and under different assumptions on the component functions, and showed that the iterates converge to

a neighborhood of the optimal solution (where the size of the neighborhood is a function of the stepsize). Most closely related to our paper is a convergence rate result provided by Nédic and Bertsekas [22], which under a strong-convexity type condition on the sum function $f(x) = \sum_{i=1}^{m} f_i(x)$, but without assuming differentiability of the component functions, shows that the distance of the iterates generated by the incremental subgradient method converges at rate $\mathcal{O}(\frac{1}{\sqrt{k}})$ to the optimal solution with a properly selected diminishing stepsize.

Luo [17] considered a special case of the problem (1.1) in dimension one when there are two convex quadratic component functions with an identical non-zero curvature and showed that IG iterates converge in this particular case at rate $\mathcal{O}(\frac{1}{k})$ to the optimal solution. Motivated by this example, in this paper we show that Nédic and Bertsekas' $\mathcal{O}(\frac{1}{\sqrt{k}})$ result can be improved when the component functions are smooth. In particular, when the component functions are quadratics and the sum function $f(x)$ is strongly convex, we first prove that the distances of the iterates generated by the IG method converge at rate $\mathcal{O}(\frac{1}{k})$ (which translates into $\mathcal{O}(\frac{1}{k^2})$ in function values by the strong convexity of f). Then, we generalize this result to twice continuously differentiable component functions under some assumptions. Achieving this rate with IG requires using a diminishing stepsize that adapts to the strong convexity constant c of the sum function, i.e., a stepsize that takes the form R/k where $R > 1/c$.[1] We then consider alternative "robust" stepsizes $\alpha_k = \Theta(\frac{1}{k^s})$ for $s \in (0, 1)$, which does not require knowledge of the strong convexity constant, and show that IG method with these stepsizes achieve a rate $\mathcal{O}(\frac{1}{k^s})$ in distances (which translates into $\mathcal{O}(\frac{1}{k^{2s}})$ in function values). We also provide lower bounds showing that these rates cannot be improved using IG. Furthermore, our results play a key role in the recently obtained convergence results for the *random reshuffling* (RR) method [14]. The random reshuffling method is a stochastic variant of IG where the order to visit the functions is selected as a random permutation of $\{1, 2, \ldots, m\}$ at the beginning of each cycle instead of the deterministic fixed order $\{1, 2, \ldots, m\}$ of IG (hence the name RR refers to the random reshuffling of the order). The convergence rate of the

[1] We note that a consequence of a paper by Hazan and Kale [15] is that when each of the component functions is strongly convex, IG with iterate averaging and stepsize $\alpha_k = R/k$ where R is the multiplicative inverse of the strong convexity constant, converges at rate $\mathcal{O}(\log k/k)$ in the suboptimality of the function value. However, the rate we obtain in this paper with a similar stepsize corresponds to $\mathcal{O}(1/k^2)$ in the suboptimality of the objective value which is much faster.

random reshuffling method has been a long standing open question. Fundamental to the analysis in [14] which solves this open problem under some technical assumptions on the objective function f is the fast convergence results introduced in this paper.

While IG achieves rate $\mathcal{O}(\frac{1}{k})$ in distances with stepsize $\frac{R}{k}$, our rate estimates show dependence on strong convexity and Lipschitz constant of the sum function. We next consider an *incremental Newton (IN) method* introduced in [13] for solving problem (1.1) which scales the gradients of the component functions with an estimate of the Hessian of the sum function $f(x)$: starting from initial point $x_1^1 \in \mathbb{R}^n$, initial stepsize $\alpha_1 > 0$ and initial Hessian estimate $H_0^1 = I_n$, for each $k \geq 1$, IN method updates the iterate x_i^k as

$$x_{i+1}^k := x_i^k - \alpha_k (\bar{H}_i^k)^{-1} \nabla f_i(x_i^k) \tag{1.4}$$

where

$$H_i^k := H_{i-1}^k + \nabla^2 f_i(x_i^k), \quad \bar{H}_i^k = H_i^k / k, \tag{1.5}$$

with the convention that $x_1^{k+1} = x_{m+1}^k$ and $H_0^{k+1} = H_m^k$. For IN, we provide rate estimates which do not depend on the Lipschitz constant and show that the IN method, unlike IG, converges with rate $\mathcal{O}(\frac{1}{k})$ without the necessity to adjust to the strong convexity constant.

Notation. For non-negative sequences $\{a_k\}$ and $\{b_k\}$, we write $a_k \geq \Omega(b_k)$ if there exists a real constant $h > 0$ and a finite integer k_0 such that $a_k \geq h b_k$ for every $k \geq k_0$. The norm $\|\cdot\|$ denotes Euclidean norm for vectors and the spectral norm for matrices. We also write $a_k \geq \tilde{\Omega}(b_k)$ if there exists a real constant $h > 0$ and infinitely many k such that $a_k \geq h b_k$ is true[2]. The matrix I_n denotes the $n \times n$ identity. The sets $\mathbb{R}_+$ and $\mathbb{N}_+$ denotes the positive real numbers and positive integers respectively. We refer to twice continuously differentiable functions on $\mathbb{R}^n$ as *smooth* functions.

2. Preliminaries. We introduce the following lemma, known as Chung's lemma, which we will make use of in our rate analysis. The proof of the part (i) of this lemma can be found in [27, Sec 2.2]. For the proof of the part (ii) of this lemma, we refer the reader to [10, Lemma 4].

[2]The $\tilde{\Omega}$ function defined here was introduced by Littlewood and Hardy in 1914. It is a weaker alternative to the Ω function and satisfies $a_k \geq \Omega(b_k) \implies a_k \geq \tilde{\Omega}(b_k)$ but not vice versa.

LEMMA 2.1. *Let $u_k \geq 0$ be a sequence of real numbers. Assume there exists k_0 such that*

$$u_{k+1} \leq (1 - \frac{a}{k^s})u_k + \frac{d}{k^{s+t}}, \quad \forall k \geq k_0,$$

where $0 < s \leq 1$, $d > 0$, $a > 0$ and $t > 0$ are given real numbers. Then,

(i) If $s = 1$, then

$$\begin{aligned} \limsup_{k\to\infty} k^t u_k &\leq \frac{d}{a-t} \qquad \text{for} \quad a > t, \\ \limsup_{k\to\infty} \frac{k^a}{\log k} u_k &< \infty \qquad \text{for} \quad a = t, \\ \limsup_{k\to\infty} k^a u_k &< \infty \qquad \text{for} \quad a < t. \end{aligned}$$

(ii) If $0 < s < 1$, then[3]

$$\limsup_{k\to\infty} k^t u_k \leq \frac{d}{a}.$$

3. Convergence rate analysis for IG.

3.1. Rate for quadratic functions. We first analyze the convergence behavior when the component functions are quadratic functions before proceeding to the more general case when functions are twice continuously differentiable. Let $f_i(x) : \mathbb{R}^n \to \mathbb{R}$ be a quadratic function of the form

$$f_i(x) = \frac{1}{2} x^T P_i x - q_i^T x + r_i, \quad i = 1, 2, \ldots, m, \tag{3.1}$$

where P_i is a symmetic $n \times n$ square matrix, $q_i \in \mathbb{R}^n$ is a column vector and r_i is a real scalar. The gradient and the Hessian of f are given by

$$\nabla f_i(x) = P_i x - q_i, \quad \nabla^2 f_i(x) = P_i. \tag{3.2}$$

The sum f is also a quadratic which we next assume to be strongly convex.

ASSUMPTION 3.1. *The sum function $f(x)$ is strongly convex on $\mathbb{R}^n$, i.e. there exists a constant $c > 0$ such that the function $f(x) - \frac{c}{2}\|x\|^2$ is convex on $\mathbb{R}^n$.*

Under this assumption, the optimal solution to the problem (1.1) is unique, which we denote by x^*. In the particular case, when each f_i is a quadratic function given by (3.2), then the Hessian matrix of the sum satisfies

$$P := \nabla^2 f(x) = \sum_{i=1}^{m} P_i \succeq cI_n \succ 0, \tag{3.3}$$

[3]Part (*ii*) of Lemma 2.1 is still correct when u_k is allowed to take negative values. However, this will not be needed in our analysis.

and the optimal solution is

$$x^* = P^{-1} \sum_{i=1}^{m} q_i. \tag{3.4}$$

The inner iterations of IG become

$$x_{i+1}^k = (I_n - \alpha_k P_i) x_i^k + \alpha_k q_i, \quad i = 1, 2, \ldots, m.$$

Therefore, the outer iterations are given by

$$\begin{aligned} x_1^{k+1} &= \prod_{i=1}^{m} (I_n - \alpha_k P_i) x_1^k + \alpha_k \sum_{i=1}^{m} \prod_{j=i+1}^{m} (I_n - \alpha_k P_j) q_i && (3.5) \\ &= \Big(I_n - \alpha_k P + \mathcal{O}(\alpha_k^3) \Big) x_1^k + \alpha_k \sum_{i=1}^{m} q_i + \alpha_k^2 T(\alpha_k) + \mathcal{O}(\alpha_k^3) && (3.6) \\ &= \big(I_n - \alpha_k P \big) x_1^k + \alpha_k \sum_{i=1}^{m} q_i + \alpha_k^2 E(\alpha_k) && (3.7) \end{aligned}$$

where

$$\begin{aligned} T(\alpha_k) &= \sum_{1 \le i < j \le m} P_j (P_i x_1^k - q_i) = \sum_{1 \le i < j \le m} P_j \nabla f_i(x_1^k) && (3.8) \\ E(\alpha_k) &= T(\alpha_k) + \mathcal{O}(\alpha_k) + \mathcal{O}(\alpha_k x_1^k). && (3.9) \end{aligned}$$

From (3.4), (3.6) and (3.7), it follows that

$$\begin{aligned} x_1^{k+1} - x^* &= \big(I_n - \alpha_k P + \mathcal{O}(\alpha_k^3) \big)(x_1^k - x^*) + \alpha_k^2 T(\alpha_k) + \mathcal{O}(\alpha_k^3). && (3.10) \\ x_1^{k+1} - x^* &= \big(I_n - \alpha_k P \big)(x_1^k - x^*) + \alpha_k^2 E(\alpha_k). \end{aligned}$$

Taking norms of both sides of the last expression, defining

$$\text{dist}_k = \|x_1^k - x^*\|$$

as the distance to the optimal solution and using the lower bound (3.3) on the eigenvalues of P, we obtain

$$\begin{aligned} \text{dist}_{k+1} &\le \big\| I_n - \alpha_k P \big\| \text{dist}_k + \alpha_k^2 \|E(\alpha_k)\| \\ &\le (1 - \alpha_k c) \text{dist}_k + \alpha_k^2 \|E(\alpha_k)\| \quad \text{if} \quad \alpha_k \|P\| \le 1. && (3.11) \\ &\le (1 - \alpha_k c) \text{dist}_k + \alpha_k^2 M_\infty \quad \text{if} \quad \alpha_k \|P\| \le 1 && (3.12) \end{aligned}$$

where

$$M_\infty := \sup_{k \ge 1} \|E(\alpha_k)\|.$$

The next theorem is based on analyzing this recursion and is on the achievable convergence rate when the component functions are quadratics.

THEOREM 3.1. *Let each* $f_i(x) = \frac{1}{2}x_i^T P_i x - q_i^T x + r_i$ *be a quadratic function as in* (3.1) *for* $i = 1, 2, \ldots, m$. *Suppose Assumption 3.1 holds. Consider the iterates* $\{x_1^k\}$ *generated by the IG method with stepsize* $\alpha_k = R/k^s$ *where* $R > 0$ *and* $s \in [0, 1]$. *Then, we have the following:*

(*i*) *If* $s = 1$, *then*

$$\begin{aligned} \limsup_{k\to\infty} k\, dist_k &\leq \frac{R^2 M}{Rc - 1} \qquad \text{for} \quad R > 1/c, \\ \limsup_{k\to\infty} \frac{k}{\log k} dist_k &< \infty \qquad \text{for} \quad R = 1/c, \\ \limsup_{k\to\infty} k^{Rc} dist_k &< \infty \qquad \text{for} \quad R < 1/c. \end{aligned}$$

where $M = \lim_{k\to\infty} \|E(\alpha_k)\| = \big\| \sum_{1\leq i<j\leq m} P_j \nabla f_i(x^*) \big\|$.

(*ii*) *If* $0 < s < 1$, *then*

$$\limsup_{k\to\infty} k^t\, dist_k \leq \frac{RM}{c}.$$

(*iii*) *If* $s = 0$ *and* $R \leq \frac{1}{\|P\|}$, *then*

$$dist_{k+1} \leq (1 - c\alpha)^k dist_1 + \frac{\alpha M_\infty}{c}, \quad \forall k \geq 1, \tag{3.13}$$

where the stepsize $\alpha = \alpha_k = R$ *is a constant and* M_∞ *is defined by* (3.12).

Proof. We first prove parts (*i*) and (*ii*). So, assume $0 < s \leq 1$. Plugging the expression for the stepsize into (3.10) and taking norms, we obtain

$$\text{dist}_{k+1} \leq \big(1 - \frac{Rc}{k^s} + \mathcal{O}(\frac{1}{k^{3s}})\big)\text{dist}_k + \frac{R^2}{k^{2s}}\|T(\alpha^k)\| + \mathcal{O}(\frac{1}{k^{3s}}). \tag{3.14}$$

It is also easy to see from (3.8) that

$$\begin{aligned} \|T(\alpha_k)\| &\leq \big\| \sum_{1\leq i<j\leq m} P_j \nabla f_i(x^*) \big\| + \big\| \sum_{1\leq i<j\leq m} P_j \big(\nabla f_i(x_1^k) - \nabla f_i(x^*)\big) \big\| \\ &= \big\| \sum_{1\leq i<j\leq m} P_j \nabla f_i(x^*) \big\| + \big\| \sum_{1\leq i<j\leq m} P_j P_i (x_1^k - x^*) \big\| \\ &\leq M + h_1 \text{dist}_k \end{aligned} \tag{3.15}$$

for some finite positive constant h_1 (that depends on $\{P_i\}_{i=1}^m$ and $\{\nabla f_i(x^*)\}_{i=1}^m$). Then, from (3.14),

$$\text{dist}_{k+1} \leq \big(1 - \frac{Rc}{k^s} + \frac{R^2 h_1}{k^{2s}} + \mathcal{O}(\frac{1}{k^{3s}})\big)\text{dist}_k + \frac{R^2 M}{k^{2s}} + \mathcal{O}(\frac{1}{k^{3s}}).$$

Finally, applying Lemma 2.1 with a choice of $0 < a < Rc$, $t = s$ and $d > R^2M$ and letting $a \to Rc$ and $d \to R^2M$ proves directly parts (i) and (ii). To prove part (iii), assume $s = 0$ and $R \leq \frac{1}{\|P\|}$. Then, the stepsize $\alpha_k = \alpha = R$ is a constant and by (3.11), for all $k \geq 1$,

$$\text{dist}_{k+1} \leq (1 - \alpha c)\text{dist}_k + \alpha^2 M_\infty.$$

From this relation, by induction we obtain for all $k \geq 1$,

$$\text{dist}_{k+1} \leq (1 - c\alpha)^k \text{dist}_1 + \alpha^2 M_\infty \sum_{j=0}^{k-1} (1 - c\alpha)^j.$$

As the geometric sum $\sum_{j=0}^{k-1}(1 - c\alpha)^j \leq \frac{1}{c\alpha}$ for all $k \geq 1$, this proves part (iii). ☐

REMARK 3.2. *Under the setting of Theorem 3.1, in the special case when x^* is a global minimizer for each of the component functions, we have $\nabla f_i(x^*) = 0$ for each $i = 1, 2, \ldots, m$. This implies that $M = 0$ and for stepsize R/k with $R > 1/c$ we have $\limsup_{k\to\infty} k\,dist_k = 0$, i.e. $dist_k = o(1/k)$. This rate can actually be improved further as follows. In this special case, assume for simplicity that $x^* = 0$ (the more general case can be treated similarly by shifting the coordinates and considering the functions $f_i(x - x^*)$ and $f(x - x^*)$). Then, this implies that $q_i = 0$ for all i and therefore from* (3.5) *we have,*

$$\begin{aligned} dist_{j+1} &= \Big\| \prod_{i=1}^{m} \big(I_n - \frac{R}{k} P_i\big) \Big\| dist_j = \Big\| I_n - \frac{R}{j} P + \mathcal{O}(1/j^2) \Big\| dist_j \\ &\leq \big(1 - \frac{Rc}{j} + \mathcal{O}(1/j^2)\big) dist_j \leq (1 - \frac{\delta}{j}) dist_j, \end{aligned}$$

where the last inequality holds for any $1 < \delta < Rc$ and j large enough. As $\prod_{j=2}^{k} \big(1 - \frac{\delta}{j}\big) \approx \prod_{j=2}^{k} \big(1 - \frac{1}{j}\big)^{\delta} = 1/k^{\delta}$, it follows that $dist_k = \mathcal{O}(1/k^{\delta})$ for any $1 < \delta < Rc$ which is a stronger statement than $dist_k = o(1/k)$.

3.2. Rate for smooth component functions. In addition to Assumption 3.1 on the strong convexity of f, we adopt the following assumptions that have appeared in a number of papers in the literature for analyzing incremental methods including [20], [2], [13].

ASSUMPTION 3.3. *The functions f_i are convex and twice continuously differentiable on $\mathbb{R}^n$ for each $i = 1, 2, \ldots, m$.*

ASSUMPTION 3.4. *The iterates $\{x_1^k, x_2^k, \ldots, x_m^k\}_{k\geq 1}$ are uniformly bounded, i.e. there exists a non-empty compact Euclidean ball $\mathcal{X} \subset \mathbb{R}^n$ that contains all the iterates.*

A consequence of these two assumptions is that by continuity of the first and second derivatives on the compact set $\mathcal{X}$, the first and second derivatives are bounded. Therefore there exists a constant G such that

$$\max_{1\le i\le m} \sup_{x\in\mathcal{X}} \|\nabla f_i(x)\| \le G \tag{3.16}$$

and there exists constants $L_i := \max_{z\in\mathcal{X}} \|\nabla^2 f_i(z)\| \ge 0$ such that

$$\|\nabla f_i(x) - \nabla f_i(y)\| \le L_i\|x-y\|, \quad \text{for all} \quad x\in\mathcal{X}, \quad i=1,2,\ldots,m. \tag{3.17}$$

From the triangle inequality, f has also Lipschitz gradients on $\mathcal{X}$ with constant

$$L = \sum_{i=1}^{m} L_i. \tag{3.18}$$

Another consequence is that, an optimal solution to the problem (1.1), which we denote by x^*, exists and is unique by the strong convexity of f. Furthermore, these two assumptions are sufficient for global convergence of both the incremental Newton and the incremental gradient methods to x^* (see [13], [4] for a more general convergence theory).

3.2.1. Analyzing IG as a gradient descent with errors. We can rewrite the inner iterations (1.2) more compactly as

$$x_1^{k+1} = x_1^k - \alpha_k\big(\nabla f(x_1^k) - e^k\big), \quad k\ge 1, \quad i=1,2,\ldots,m, \tag{3.19}$$

where the term

$$e^k = \sum_{i=1}^{m}\big(\nabla f_i(x_1^k) - \nabla f_i(x_i^k)\big) \tag{3.20}$$

can be viewed as the gradient error. If Assumption 3.3 holds, we can substitute

$$\nabla f(x_1^k) = A_k(x_1^k - x^*)$$

into (3.19) where $A_k = \displaystyle\int_0^1 \nabla^2 f(x^* + \tau(x^k - x^*))d\tau$ is an average of the Hessian matrices on the line segment $[x_1^k, x^*]$ to obtain

$$x_1^{k+1} - x^* = (I_n - \alpha_k A_k)(x_1^k - x^*) + \alpha_k e^k, \quad k\ge 1, \quad i=1,2,\ldots,m. \tag{3.21}$$

Taking norms of both sides, this implies that

$$\text{dist}_{k+1} \le \|I_n - \alpha_k A_k\|\text{dist}_k + \alpha_k\|e^k\|. \tag{3.22}$$

These relations show that the evolution of the distance to the optimal solution is controlled by the decay of the stepsize α_k and the gradient error $\|e^k\|$. This motivates deriving tight upper bounds for the gradient error. Note also that under Assumptions 3.1, 3.3 and 3.4, the Hessian of f and the averaged Hessian matrix A_k admit the bounds

$$cI_n \preceq \nabla^2 f(x), A_k \preceq LI_n, \quad x \in \mathcal{X}. \tag{3.23}$$

(see also (3.18)). The gradient error consists of the difference of gradients evaluated at different inner steps (see (3.20)). This error can be controlled by the Lipschitzness of the gradients as follows: For any $k \geq 1$,

$$\begin{aligned}\|e^k\| &\leq \sum_{i=2}^{m} L_i \|x_1^k - x_i^k\| \leq \sum_{i=2}^{m} L_i \sum_{j=1}^{i-1} \|x_j^k - x_{j+1}^k\| \\ &\leq \sum_{i=2}^{m} L_i \alpha_k \sum_{j=1}^{i-1} \|\nabla f_j(x_j^k)\| \\ &\leq \alpha_k \, \widetilde{M}, \end{aligned} \tag{3.24}$$

where

$$\widetilde{M} := LGm, \tag{3.25}$$

L is a Lipschitz constant for the gradient of f as in (3.18) and G is an upper bound on the gradients as in (3.16). Finally, plugging this into (3.22) and using the bounds (3.23) on the eigenvalues of A_k,

$$\begin{aligned}\text{dist}_{k+1} &\leq \max(\|1 - \alpha_k c\|, \|1 - \alpha_k L\|)\text{dist}_k + \alpha_k^2 \, \widetilde{M} \\ &\leq (1 - \alpha_k c)\text{dist}_k + \alpha_k^2 \, \widetilde{M} \quad \text{if} \quad \alpha_k L \leq 1. \end{aligned} \tag{3.26}$$

This is the analogue of the recursion (3.12) obtained for quadratics with the only difference that the constants M_∞ and $\|P\|$ are replaced by their analogues $\widetilde{M}$ and L respectively. Then, a reasoning along the lines of the proof of Theorem 3.1 yields the following convergence result which generalizes Theorem 3.1 from quadratic functions to smooth functions just by modifying the constants properly. We skip the proof for the sake of brevity.

THEOREM 3.2. *Let $f_i(x) : \mathbb{R}^n \to \mathbb{R}$, $i = 1, 2, \ldots, m$ be component functions satisfying Assumptions 3.1 and 3.3. Consider the iterates $\{x_1^k, x_2^k, \ldots, x_m^k\}_{k \geq 1}$ obtained*

by the IG iterations (1.2) *with a decaying stepsize* $\alpha_k = R/k^s$ *where* $R > 0$ *and* $s \in [0, 1]$. *Suppose that Assumption 3.4 is also satisfied. Then, all the conclusions in parts* (i), (ii) *and* (iii) *of Theorem 3.1 remain valid by replacing* M *and* M_∞ *with* $\widetilde{M}$ *and replacing* $\|P\|$ *with* L.

REMARK 3.5. *Under the conditions of Theorem 3.1, the quadratic functions* f_i *have Lipschitz continuous gradients with constants* $L_i = \|P_i\|$. *Therefore,*

$$M \leq \sum_{1 \leq i \leq m} \sum_{j=i+1}^{m} L_j \, \|\nabla f_i(x^*)\| \leq \sum_{1 \leq i \leq m} L \, \|\nabla f_i(x^*)\| \leq \widetilde{M}$$

by the definitions of L *and* $\widetilde{M}$ *from* (3.18) *and* (3.25). *This shows that the rate results obtained in Theorem 3.1 with decaying stepsize for the special case of quadratic functions is tighter than that of the general case of smooth functions obtained in Theorem 3.2 as expected.*

Under a strong convexity-type condition and subgradient boundedness, Nédic and Bertsekas consider the IG method with constant stepsize and show that when f_i are convex but not necessarily smooth or differentiable, for any given $\varepsilon > 0$, it suffices to have $\mathcal{O}(\log(\frac{1}{\varepsilon})/\varepsilon^2)$ cycles of IG for converging to the ε-neighborhood $\{x \in \mathbb{R}^n : \|x - x^*\| \leq \varepsilon\}$ of an optimal solution x^* [22, Proposition 2.4]. The following Corollary of Theorem 3.2 shows that this result can be improved to $\mathcal{O}(\log(\frac{1}{\varepsilon})/\varepsilon)$ when the component functions are smooth and strongly convex.

COROLLARY 3.3. *Under the conditions of Theorem 3.2, let* $\varepsilon < 2 \, \widetilde{M} \, /c^2$ *be a given positive number. IG with constant stepsize* $\alpha = \varepsilon c/(2 \, \widetilde{M})$ *requires at most*

$$\mathcal{O}\left(\frac{\widetilde{M}}{c^2} \frac{\log(1/\varepsilon)}{\varepsilon}\right) \tag{3.27}$$

cycles to guarantee convergence to an ε*-neighborhood of the optimal solution* x^*.

Proof. Given such $\varepsilon > 0$ and stepsize α, we note that $c\alpha < 1$ and $\alpha \, \widetilde{M}/c = \varepsilon/2$. Furthermore, by Theorem 3.2, the inequality (3.13) holds with M_∞ replaced by $\widetilde{M}$. Therefore, there exists a constant K such that

$$(1 - c\alpha)^k \text{dist}_1 \leq \exp(-c\alpha k)\text{dist}_1 < \frac{\varepsilon}{2}, \quad \forall k \geq K,$$

and $\text{dist}_{k+1} < \varepsilon$ for all $k \geq K$, i.e. the iterates lie inside an ε-neighborhood of the optimizer after K cycles. By taking log of both sides and using $\log(1 - z) \approx z$ for z around zero, straightforward calculations show that this condition is satisfied for K satisfying (3.27). ∎

The rate results in Theorem 3.1 for $\mathcal{O}(1/k)$ stepsize (when $s = 1$) requires adjusting the stepsize parameter R to the strong convexity parameter c which requires the estimation or the knowledge of a lower bound for c (need to choose $R > 1/c$). The following example illustrates that the convergence can be slow when R is not properly adjusted to c. Similar issues with $1/k$-decay step sizes are also widely noted in the analysis of the stochastic gradient descent method in the stochastic approximation literature, see e.g. [1, 10, 21, 25].

EXAMPLE 3.6. *Let $f_i(x) = x^2/20$ for $i = 1, 2$, $x \in \mathbb{R}$. Then, we have $m = 2$, $c = 1/5$ and $x^* = 0$. Take $R = 1$ which corresponds to the stepsize $1/k$. The IG iterations are*

$$x_1^{k+1} = \left(1 - \frac{1}{10k}\right)^2 x_1^k.$$

If $x_1 = 1$, a simple analysis similar to [25] shows $x_1^k = dist_k > \Omega(\frac{1}{k^{1/5}})$.

3.3. Lower bounds. Consider the following set of quadratic functions which are stronly convex with parameter c and have Lipschitz gradients with constant L:

$$\mathcal{C}_{c,L} = \bigcup_{n=1}^{\infty} \left\{ \bar{f}(x) = \frac{1}{2} x^T P x - q^T x + r \;\middle|\; P \text{ symmetric}, cI_n \preceq P \preceq LI_n; x, q \in \mathbb{R}^n; r \in \mathbb{R} \right\}.$$

Theorem 3.1 and Remark 3.5 shows that when IG is applied to quadratic functions $\bar{f}_i : \mathbb{R}^n \to \mathbb{R}$ with a sum $\bar{f} \in \mathcal{C}_{c,L}$ using a stepsize $\alpha_k = R/k$ where R is properly chosen as a function of the strong convexity constant c, it results in

$$\limsup_{k\to\infty} k\text{dist}_k \leq \frac{M}{Rc - 1} \leq \frac{\widetilde{M}}{Rc - 1} = \frac{LGm}{Rc - 1}.$$

In other words, $\text{dist}_k = \mathcal{O}(1/k)$. A natural question is whether one could improve this rate by choosing a perhaps alternative stepsize. Would it be possible to obtain $\text{dist}_k = o(1/k)$ uniformly (for every such $\{\bar{f}_i\}_{i=1}^m$ and m)? The next result gives a negative answer to this question showing that no matter which stepsize we choose, there exists simple quadratic functions which result in iterates $\{x_1^k\}$ satisfying $\text{dist}_k \geq \tilde{\Omega}(1/k)$.

THEOREM 3.4. *Consider the following IG iterations applied to quadratic component functions $\bar{f}_i : \mathbb{R}^n \to \mathbb{R}$ where $\bar{f} = \left(\sum_{i=1}^m \bar{f}_i\right) \in \mathcal{C}_{c,L}$:*

$$x_{i+1}^k = x_i^k - \sigma(c, L, k)\nabla \bar{f}_i(x_i^k), \quad k \geq 1, \quad i = 1, 2, \ldots, m,$$

where the stepsize sequence $\alpha_k = \sigma(c, L, k) : \mathbb{R}^3_+ \to \mathbb{R}_+$ *is independent from the choice of each* $\bar{f}_i$*. Suppose that for every choice of* m, n *and such* $\{\bar{f}_i : \mathbb{R}^n \to \mathbb{R}\}_{i=1}^m$*, we have*

$$\limsup_{k\to\infty} k\, dist_k \leq \bar{b}$$

where $\bar{b} > 0$ *depends only on* L, G, m, c *and* σ*. Then, the following statements are true:*

1. *The stepsize sequence satisfies* $\limsup_{k\to\infty} k\alpha_k \geq \underline{b}$ *where* $\underline{b} = \bar{b}/L$*.*
2. *There exists positive integers* $\tilde{m}$*,* $\tilde{n}$ *and functions* $\{\tilde{f}_i : \mathbb{R}^{\tilde{n}} \to \mathbb{R}\}_{i=1}^{\tilde{m}}$ *such that* $\tilde{f} = \left(\sum_{i=1}^{\tilde{m}} \tilde{f}_i\right) \in \mathcal{C}_{c,L}$ *and the iterates* $\{x_1^k\}$ *generated by the IG applied to* $\tilde{f} = \sum_{i=1}^m \tilde{f}_i$ *satisfy*

$$dist_k \geq \tilde{\Omega}(1/k).$$

Proof.

(i) We follow the approach of [28, Appendix A] which was introduced to justify the optimality of the $\Theta(1/k)$ stepsize for the stochastic gradient descent method. Consider the simple example $\bar{f}(x) = \bar{f}_1(x) = \frac{L}{2}x^2 \in \mathcal{C}_{c,L}$ with only one component function in dimension one ($m = n = 1$) or a similar alternative example $\bar{f}(x) = \frac{L}{2}\big(x(1)^2 + x(2)^2\big) \in \mathcal{C}_{c,L}$ with two component functions $\bar{f}_1(x) = \frac{L}{2}x(1)^2$ and $\bar{f}_2(x) = \frac{L}{2}x(2)^2$ in dimension two ($n = m = 2$) where $x(\ell)$ denotes the ℓ-th coordinate of the vector x. In any of these two examples, IG becomes the classical gradient descent method leading to the iterations $x_1^{k+1} = \prod_{j=1}^k (1 - \alpha_j L)x_1^1$. By the assumption, we have at least

$$\left|\prod_{j=1}^k (1 - \alpha_j L)\right| \leq \frac{\bar{b}}{k} + o(\frac{1}{k}) \leq \frac{2\bar{b}}{k} \quad \text{for } k \text{ large} \tag{3.28}$$

and $\alpha_k \to 0$ (otherwise simple examples show the global convergence may not happen from an arbitrary initial point). By taking the natural logarithm of both sides, this is equivalent to requiring

$$\sum_{j=1}^k -\ln|1 - \alpha^j L| \geq 2\bar{b}\log k \quad \text{for } k \text{ large.}$$

Using $2z \geq -\ln(1 - z)$ for $0 \leq z \leq \frac{1}{2}$, it follows that

$$\sum_{j=1}^k \alpha_j \geq \underline{b}\log k \tag{3.29}$$

when k is large enough. Assume there exists δ such that $\limsup_{k\to\infty} k\alpha_k < \delta < \underline{b}$. Then, by definition of the limit superior, we have $\alpha_k \leq \frac{\delta}{k}$ for any k large enough. By summing this inequality over the iterations k, we obtain $\sum_{j=1}^{k} \alpha_j \leq \delta \log(k) + b_2$ for a constant b_2 and for any k large enough. This contradicts with (3.29). Therefore, no such δ exists, i.e. $\limsup_{k\to\infty} k\alpha_k \geq \underline{b}$. This completes the proof.

(*ii*) Consider the following simple example with two quadratics $\bar{f} = \bar{f}_1 + \bar{f}_2$ with $\bar{f}_1(x) = \frac{L}{2}(x-1)^2$ and $\bar{f}_2(x) = \frac{L}{2}(x+1)^2$ in dimension one ($m = 2$, $n = 1$). Then, applying IG with an initial point $x_1^1 \in \mathbb{R}$ results in the iterates $\{x_1^k, x_2^k\}$ with

$$x_2^k = x_1^k - \bar{\alpha}_k(x_1^k - 1), \tag{3.30}$$

$$x_1^{k+1} = (1-\bar{\alpha}_k)^2 x_1^k - (\bar{\alpha}_k)^2, \tag{3.31}$$

$$x_2^{k+1} = (1-\bar{\alpha}_k)^2 x_2^k + (\bar{\alpha}_k)^2, \tag{3.32}$$

where $\bar{\alpha}_k = \alpha_k L$ is the normalized stepsize. Define $y^k = x_1^k + x_2^k$. By summing up (3.31) and (3.32), we see that

$$y^{k+1} = (1-\bar{\alpha}_k)^2 y^k = \prod_{j=1}^{k} (1-\bar{\alpha}_j)^2 y^1. \tag{3.33}$$

By the necessary condition (3.28), we also have

$$0 \leq |y^k| \leq \mathcal{O}(1/k^2). \tag{3.34}$$

Finally, plugging $y^k = x_1^k + x_2^k$ into (3.30), we obtain

$$x_1^k = \frac{y^k}{2} + \bar{\alpha}_k \frac{(x_1^k - 1)}{2}.$$

As $\alpha_k = \tilde{\Omega}(1/k)$ by part (*i*) and x_1^k is converging to zero, it follows from (3.34) and the triangle inequality that

$$|x_1^k| = \text{dist}_k \geq \bar{\alpha}_k \frac{|x_1^k - 1|}{2} - \frac{|y^k|}{2} = \tilde{\Omega}(1/k).$$

This completes the proof.

□

3.3.1. Lower bounds for stepsize $\alpha_k = R/k^s$ with $s \in [0,1]$. In this section, we are interested in the number of cycles necessary with IG to reach to an

ε-neighborhood of an optimal solution using the stepsize $\alpha_k = R/k^s$ with $s \in [0,1]$. As before, we consider the case when the component functions are smooth and their sum is strongly convex.

When $s = 0$, the stepsize $\alpha_k = \alpha$ is a constant and there exists simple examples (with two quadratics in dimension one) which necessitate $\Omega\big(\log(1/\varepsilon)/\varepsilon\big)$ cycles to reach to an ε-neighborhood of an optimal solution (see [17, Proposition 2.2]). Therefore, it can be argued that the dependancy on ε of the iteration complexity in Corollary 3.3 cannot be improved further.

For the special case $s \in (1/2, 1]$, Luo gives an analysis which suggests that one would expect to have $\text{dist}_k = \Omega(1/k^s)$ for least square problems (see [17, Remark 2, after the proof of Theorem 3.1]). We next provide a rigorous lower bound for the more general case $s \in (0,1)$. Specifically, the simple example given in the proof of part (ii) of Theorem 3.4 provides a lower bound for $s \in (0,1]$ by an analysis almost identical to the proof of part (ii) of Theorem 3.4 leading to the following result:

LEMMA 3.5. *Consider the iterates $\{x_1^k\}$ generated by the IG method with decaying stepsize $\alpha = R/k^s$ where $s \in (0,1]$ applied to quadratic component functions f_i whose sum f is strongly convex. Then, $dist_k = \Omega(1/k^s)$.*

This lower bound is based on an example in dimension one. However, one can also construct similar examples in higher dimensions. In Appendix B, we provide an alternative example in dimension two for illustrating this fact, although it is not as simple as the dimension one example.

4. Convergence rate analysis for IN. To analyze the gradient errors introduced in the IN iterations (1.4)–(1.5), we rewrite the outer IN iterations using [13, Equation 2.12] as:

$$x_1^{k+1} = x_1^k - \alpha_k (\bar{H}_m^k)^{-1} (\nabla f(x_1^k) + e_g^k) \tag{4.1}$$

where

$$e_g^k = \sum_{j=1}^{m} \left(\nabla f_j(x_j^k) - \nabla f_j(x_1^k) + \frac{1}{\alpha_k k} \nabla^2 f_j(x_j^k)(x_1^k - x_j^k) \right) \tag{4.2}$$

is the gradient error and

$$\bar{H}_m^k = \frac{H_0^1 + \sum_{i=1}^{k} \sum_{j=1}^{m} \nabla^2 f_j(x_j^i)}{k} = \frac{\sum_{i=1}^{k} \nabla^2 f(x_1^i)}{k} + e_h^k \tag{4.3}$$

is an averaged Hessian up to an error term

$$e_h^k = \frac{H_0^1 + \sum_{i=1}^k \sum_{j=1}^m \left(\nabla^2 f_j(x_j^i) - \nabla^2 f_j(x_1^i) \right)}{k}. \tag{4.4}$$

We let $\alpha_k = R/k$ and introduce the following norm:

$$\|z\|_* := \left(z^T H_* z\right)^{1/2}, \quad z \in \mathbb{R}^n \quad \text{where} \quad H_* := \nabla^2 f(x^*). \tag{4.5}$$

which arises in the analysis of the self-concordant functions and Newton's method [26]. The next theorem shows that unlike IG, IN can achieve the $\mathcal{O}(1/k)$ rate without requiring to know or estimate the strong convexity constant of f. Furthermore, the constants arising in IN when considered in the $*$-norm do not depend on the Lipschitz constant L unlike IG.

THEOREM 4.1. *Let f_i be component functions satisfying Assumptions 3.1 and 3.3. Consider the the iterates $\{x_1^k, \ldots, x_m^k\}$ generated by the IN method with stepsize $\alpha_k = R/k$ where $R > 1$. Assume that the iterates are uniformly bounded. Then, we have*

$$\limsup_{k \to \infty} k \|x_1^k - x_*\|_* \leq \frac{BR(R+1)}{R-1} \tag{4.6}$$

where $\|\cdot\|_*$ and H_* are defined by (4.5) *and $B = \sum_{i=1}^m \|H_*^{-1/2} \nabla f_i(x^*)\| \leq G/\sqrt{c}$ where G is defined by* (3.16).

The proof of this Theorem is given in the Appendix A. The main idea is to change variables $y = H_*^{1/2} x$ and analyze the corresponding iterates $y_1^k = H_*^{1/2} x_1^k$. By this change of variables, it can be shown that $\{y_1^k\}$ follows a similar recursion to the IN iterates $\{x_1^k\}$ converging to $y_* = H_*^{1/2} x_*$. Then, one can analyze how fast the sequence $\|y_1^k - y_*\| = \|x_1^k - x_*\|_*$ decays to zero by exploiting the fact that y-coordinates have the advantage that the local strong convexity constant and the local Lipschitz constant of f around y_* are both equal to one due to the normalization obtained by this change of variable.

5. Conclusion. We analyzed the convergence rate of the IG and IN algorithms when the component functions are smooth and the sum of the component functions is strongly convex. This covers the interesting case of many regression problems including the ℓ_2 regularized regression problems. For IG, we show that the distance of the iterates converges at rate $\mathcal{O}(1/k^s)$ to the optimal solution with a diminishing

stepsize of the form $\alpha_k = \mathcal{O}(1/k^s)$ for $s \in (0, 1]$. This improves the previously known $\mathcal{O}(1/\sqrt{k})$ rate (when $s \in (1/2, 1]$) and translates into convergence at rate $\mathcal{O}(1/k^{2s})$ of the suboptimality of the objective value. For constant stepsize, we also improve the existing iteration complexity results for IG from $\mathcal{O}(\frac{\log(1/\varepsilon)}{\varepsilon^2})$ to $\mathcal{O}(\frac{\log(1/\varepsilon)}{\varepsilon})$ to reach to an ε-neighborhood of an optimal solution. In addition, we provided lower bounds showing that these rates cannot be improved using IG.

Achieving the fastest $\mathcal{O}(1/k)$ rate in distances with IG requires a good knowledge or approximation of the strong convexity constant of the sum function f. However, we showed that IN as a second-order method, can achieve this fast rate without the knowledge of the strong convexity constant.

Appendix A. Proof of Theorem 4.1.

Proof. By a change of variable let $y = H_*^{1/2}x$ and define $\hat{f}(y) = f(x)$. Consider the IN iterates in the y-coordinates. By the chain rule, we have

$$\nabla f(x) = H_*^{1/2}\nabla \hat{f}(y), \quad \nabla^2 f(x) = H_*^{1/2}\nabla^2 \hat{f}(y)H_*^{1/2}. \tag{A.1}$$

Using these identities, the IN iterations (1.4)–(1.5) become

$$y_{i+1}^k := y_i^k - \alpha_k(\bar{D}_i^k)^{-1}\nabla \hat{f}_i(y_i^k), \quad i = 1, 2, \ldots, m, \tag{A.2}$$

where $\bar{D}_i^k = D_i^k/k$ with

$$y_i^k = H_*^{1/2}x_i^k, \quad D_i^k := D_{i-1}^k + \nabla^2 \hat{f}_i(y_i^k) = H_*^{-1/2}H_i^k H_*^{-1/2}. \tag{A.3}$$

Furthermore, the IN method is globally convergent under these assumptions (see [13]), i.e. $x_1^k \to x^*$. More generally, due to the cyclic structure, we have also $x_i^k \to x^*$ for each $i = 1, 2, \ldots, m$. Then, from the Hessian update formula (1.5), it follows that $\bar{H}_i^k \to H_*$ for each $i = 1, 2, \ldots, m$ fixed and

$$\bar{D}_i^k \to \nabla^2 \hat{f}(y^*) = H_*^{-1/2}H_*H_*^{-1/2} = I_n, \quad i = 1, 2, \ldots, m, \tag{A.4}$$

where we used the second change of variable identity from (A.1) to calculate $\nabla^2 \hat{f}(y^*)$. Comparing the IN iterations (1.4) in the x-coordinates and the IN iterations (A.2) in the y-coordinates, we see that they have exactly the same form, the only differences are that in the latter the gradients and the Hessian matrices are taken with respect to y (instead of x) and f is replaced with $\hat{f}$. Therefore, the inequalities (4.1)–(4.2) hold if we replace f with $\hat{f}$ and x_j^i with y_j^i leading to

$$y_1^{k+1} = y_1^k - \alpha_k(\bar{D}_k)^{-1}(\nabla f(y_1^k) + e_y^k), \quad \text{where} \quad \bar{D}_k := \bar{D}_m^k, \tag{A.5}$$

and the gradient error becomes

$$e_y^k = \sum_{j=1}^{m} \left(\nabla \hat{f}_j(y_j^k) - \nabla \hat{f}_j(y_1^k) + \frac{1}{R} \nabla^2 \hat{f}_j(y_j^k)(y_1^k - y_j^k) \right) \tag{A.6}$$

where we set $\alpha_k = R/k$. Setting $\nabla \hat{f}(y_1^k) = Y_k(y_1^k - y^*)$ in (A.5) with an averaged Hessian

$$Y_k = \int_0^1 \nabla^2 \hat{f}(y^* + \tau(y_1^k - y^*)) d\tau \tag{A.7}$$

where $y^* = H_*^{1/2} x_*$ and using the triangle inequality we obtain

$$\|y_1^{k+1} - y^*\| \leq \underbrace{\left\| \left(I_n - \frac{R}{k} \bar{D}_k^{-1} Y_k\right)(y_1^k - y^*) \right\|}_{:= m_k} + \frac{R}{k} \underbrace{\|(\bar{D}_k)^{-1} e_y^k\|}_{:= n_k}. \tag{A.8}$$

The remaining of the proof consists of estimating the terms m_k and n_k on the right-hand side separately in the following three steps which gives the desired convergence rate of the left-hand side $\|y_1^{k+1} - y^*\| = \|x_1^{k+1} - x^*\|_*$:

Step 1: (Bounding m_k) We first observe that

$$m_k^2 = \left\| \left(I_n - \frac{R}{k} \bar{D}_k^{-1} Y_k\right)(y_1^k - y^*) \right\|^2 = (y_1^k - y^*)^T S_k (y_1^k - y^*) \tag{A.9}$$

where

$$S_k = I_n - \frac{R}{k} Z_k, \quad Z_k = Y_k \bar{D}_k^{-1} + \bar{D}_k^{-1} Y_k - \frac{R}{k} Y_k \bar{D}_k^{-2} Y_k. \tag{A.10}$$

From (A.4), we have $\bar{D}_k = \bar{D}_m^k \to I_n$. Furthermore, as y_1^k converges to y^* by the global convergence of IN, Y_k defined in (A.7) converges to $\nabla_y^2 f(y^*) = I_n$ as well. Therefore, we have $Z_k = 2I_n + o(1)$ in (A.10) which leads to

$$S_k = (1 - \frac{2R}{k}) I_n + o(\frac{1}{k}).$$

Then, for every $\varepsilon \in (0, 1)$, there exists a finite $k_1 = k_1(\varepsilon)$ such that for $k \geq k_1(\varepsilon)$,

$$S_k \preceq \left(1 - \frac{2R(1-\varepsilon)}{k}\right) I_n$$

and therefore

$$m_k^2 = (y_1^k - y^*)^T S_k (y_1^k - y^*) \leq \left(1 - \frac{2R(1-\varepsilon)}{k}\right) \|y_1^k - y^*\|^2$$

for $k \geq k_1(\varepsilon)$. By taking the square roots of both sides, for $k \geq \max\{k_1, 2R\}$, we obtain

$$m_k \leq \left(1 - \frac{R(1-\varepsilon)}{k}\right)\|y_1^k - y^*\| \tag{A.11}$$

where we used $(1-z)^{1/2} \leq 1 - z/2$ for $z \in [0,1]$ with $z = \frac{2R(1-\varepsilon)}{k}$.

Step 2: (Bounding n_k) Similarly we can write

$$\nabla \hat{f}_j(y_j^k) - \nabla \hat{f}_j(y_1^k) = Y_{k,j}(y_j^k - y_1^k)$$

with an averaged Hessian

$$Y_{k,j} = \int_0^1 \nabla^2 \hat{f}_j(y_1^k + \tau(y_j^k - y_1^k))d\tau \underset{k\to\infty}{\to} \nabla^2 \hat{f}_j(y^*) \preceq \sum_{i=1}^m \nabla^2 \hat{f}_i(y^*) = \nabla^2 \hat{f}(y^*) = I_n \tag{A.12}$$

as $k \to \infty$ for all $j = 1, 2, \ldots, m$ where we used (A.4) in the last equality and the fact that $\nabla^2 f_i(y^*) \succeq 0$ implied by the convexity of f_i. Next, we decompose the gradient error term (A.6) into two parts as

$$e_y^k = e_{y,1}^k + e_{y,2}^k$$

with

$$e_{y,1}^k = \sum_{j=1}^m Y_{k,j}(y_j^k - y_1^k), \quad e_{y,2}^k = \frac{1}{R}\sum_{j=1}^m \nabla^2 \hat{f}_j(y_j^k)(y_1^k - y_j^k).$$

From the triangle inequality for n_k defined in (A.8), we have

$$n_k \leq \sum_{\ell=1}^2 n_{k,\ell} \quad \text{with} \quad n_{k,\ell} = \|\bar{D}_k^{-1} e_{y,\ell}^k\|. \tag{A.13}$$

We then estimate $n_{k,\ell}$ for $\ell = 1$ and $\ell = 2$:

$$\begin{aligned} n_{k,1} &= \|\sum_{j=1}^m \bar{D}_k^{-1} Y_{k,j}(y_j^k - y_1^k)\| \\ &= \frac{R}{k}\left\|\sum_{j=1}^m \sum_{\ell=1}^{j-1} \bar{D}_k^{-1} Y_{k,j}\big(\bar{D}_\ell^k\big)^{-1} \nabla \hat{f}_\ell(y_\ell^k)\right\|. \end{aligned} \tag{A.14}$$

From (A.4) and (A.12), for every $\ell, j \in \{1, 2, \ldots, m\}$, each summand above in the last equality satisfies

$$\lim_{k\to\infty} \bar{D}_k^{-1} Y_{k,j}\big(\bar{D}_\ell^k\big)^{-1} \nabla \hat{f}_\ell(y_\ell^k) = \nabla^2 \hat{f}_j(y^*) \nabla \hat{f}_\ell(y^*) \tag{A.15}$$

so that

$$\begin{aligned}\lim_{k\to\infty} kn_{k,1} &= R\Big\|\sum_{\ell=1}^{m}\sum_{j=\ell+1}^{m}\nabla^2\hat{f}_j(y^*)\nabla\hat{f}_\ell(y^*)\Big\| \\ &\le R\sum_{\ell=1}^{m}\Big\|\sum_{j=\ell+1}^{m}\nabla^2\hat{f}_j(y^*)\Big\|\|\nabla\hat{f}_\ell(y^*)\| \le R\sum_{\ell=1}^{m}\|\nabla^2\hat{f}(y^*)\|\|\nabla\hat{f}_\ell(y^*)\| \\ &\le RB\end{aligned}$$

where in the last step we used the fact that $\nabla^2\hat{f}(y^*) = I_n$ and the change of variable formula (A.1) on gradients. Similarly,

$$\begin{aligned}n_{k,2} = \|\bar{D}_k^{-1}e_{y,2}^k\| &= \frac{1}{R}\Big\|\sum_{j=1}^{m}\bar{D}_k^{-1}\nabla^2\hat{f}_j(y_j^k)(y_j^k - y_1^k)\Big\| \\ &= \frac{1}{k}\Big\|\sum_{j=1}^{m}\bar{D}_k^{-1}\nabla^2\hat{f}_j(y_j^k)\sum_{\ell=1}^{j-1}(\bar{D}_\ell^k)^{-1}\nabla\hat{f}_\ell(y_\ell^k)\Big\|\end{aligned}$$

Then, as $\nabla^2\hat{f}_j(y_j^k) \to \nabla^2\hat{f}_j(y^*)$, it follows similarly from (A.4) that

$$\lim_{k\to\infty} kn_{k,2} = \Big\|\sum_{j=1}^{m}\sum_{\ell=1}^{j-1}\nabla^2\hat{f}(y^*)\nabla\hat{f}(y^*)\Big\| \le B. \tag{A.16}$$

Going back to the triangle inequality bound (A.13) on n_k, we arrive at

$$\begin{aligned}\lim_{k\to\infty} kn_k &\le \limsup_{k\to\infty} kn_{k,1} + \limsup_{k\to\infty} kn_{k,2} \\ &\le RB + B = (R+1)B.\end{aligned}$$

In other words, for any $\varepsilon > 0$, there exists $k_2 = k_2(\varepsilon)$ such that

$$n_k \le (1+\varepsilon)(R+1)B\frac{1}{k}, \quad \forall k \ge k_2(\varepsilon). \tag{A.17}$$

Step 3: (Deriving the rate) Let $\varepsilon \in (0, \frac{R-1}{2R})$ so that $R_\varepsilon := R(1-\varepsilon) > 1$. Then, it follows from (A.8), (A.11) and (A.17) that for $k \ge \max\{k_1(\varepsilon), 2R, k_2(\varepsilon)\}$,

$$\|y_1^{k+1} - y^*\| \le \big(1 - \frac{R_\varepsilon}{k}\big)\|y_1^k - y^*\| + \frac{(1+\varepsilon)BR(R+1)}{k^2}.$$

Applying Lemma 2.1 with $u_k = \|y_1^k - y^*\| = \|x_1^k - x^*\|_*$, $a = R_\varepsilon > 1$ and $s = 1$ leads to

$$\limsup_{k\to\infty} k\|x_1^k - x_*\|_* \le \frac{(1+\varepsilon)BR(R+1)}{R(1-\varepsilon)-1}. \tag{A.18}$$

Letting $\varepsilon \to 0$, completes the proof.

□

Appendix B. An example in dimension two with $\text{dist}_k = \Omega(1/k^s)$. The aim is to construct a set of component functions in dimension two such that if IG is applied with stepsize $\alpha_k = \Theta(1/k^s)$ with $0 < s \leq 1$, the resulting iterates satisfy $\text{dist}_k \geq \Omega(1/k^s)$.

Consider the following least squares example in dimension two ($n = 2$) with $m = 8$ quadratics defined as

$$\tilde{f}_i(x) = \frac{1}{2}(c_i^T x + 1)^2, \quad i = 1, 2, \ldots, 8,$$

where the vectors $c_i \in \mathbb{R}^2$ are

$$c_1 = c_6 = -c_2 = -c_5 = [-1, 0]^T, \tag{B.1}$$

$$c_3 = c_8 = -c_4 = -c_7 = [0, -1]^T. \tag{B.2}$$

It is easy to check that the sum $\tilde{f} := \sum_{i=1}^8 \tilde{f}_i$ is strongly convex as

$$\nabla^2 \tilde{f}(x) = \sum_{i=1}^8 c_i c_i^T = 4I_n \succ 0. \tag{B.3}$$

Starting from an initial point $\tilde{x}_1^1$, the IG method with stepsize α_k leads to the iterations

$$\tilde{x}_{i+1}^k = (I_n - \alpha_k c_i c_i^T)\tilde{x}_i^k - \alpha_k c_i, \quad i = 1, 2, \ldots, 8, \tag{B.4}$$

which implies

$$\begin{aligned} \tilde{x}_1^{k+1} &= \prod_{i=1}^8 (I_n - \alpha_k c_i c_i^T)\tilde{x}_1^k - \alpha_k \sum_{i=1}^8 c_i + \alpha_k^2 \sum_{1 \leq i < j \leq 8} (c_j^T c_i) c_j + \mathcal{O}(\alpha_k^3) \\ \tilde{x}_1^{k+1} &= \prod_{i=1}^8 (I_n - \alpha_k c_i c_i^T)\tilde{x}_1^k + \mathcal{O}(\alpha_k^3) \end{aligned} \tag{B.5}$$

where in the second step we used the fact that the terms with α_k and α_k^2 above vanish due to symmetry properties imposed by relations (B.1) and (B.2). The cyclic order $\{1, 2, \ldots, 8\}$ is special in the sense that it takes advantage of the symmetry in the problem leading to cancellations of the $\mathcal{O}(\alpha_k)$ and $\mathcal{O}\big(\alpha_k^2\big)$ terms leading to smaller $\mathcal{O}\big(\alpha_k^3\big)$ additive error terms, whereas it can be checked that this is not the case for the order $\{2, 3, \ldots, 8, 1\}$. With this intiution in mind, we next show that the sequence $\{\tilde{x}_2^k\}$ converges to the optimal solution $x^* = 0$ slower than the sequence $\{\tilde{x}_1^k\}$ does.

Using (B.3), the fact that $x^* = 0$ for this specific example and the triangle inequality on (B.5),

$$\begin{aligned}\text{dist}_{k+1} &\leq \|\prod_{i=1}^{8}\big(I_n - \alpha_k c_i c_i^T\big)\|\text{dist}_k + \mathcal{O}(\alpha_k^3)\\ &\leq \big|1 - 4\alpha_k + \mathcal{O}(\alpha_k^2)\big|\text{dist}_k + h_3(\alpha_k)^3\end{aligned}$$

for some constant $h_3 > 0$. As $\alpha_k = \Theta(1/k^s)$, applying part (ii) of Lemma (2.1) with $t = 2s$ gives

$$\|\tilde{x}_1^k\| = \mathcal{O}(1/k^{2s}). \tag{B.6}$$

Then, for $i = 1$ the inner iterations (B.4) gives

$$\tilde{x}_2^k = \tilde{x}_1^k - \alpha_k(c_1^T\tilde{x}_1^k + 1)c_1.$$

As $\tilde{x}_1^k \to 0$, $(c_1^T\tilde{x}_1^k + 1)c_1 \to c_1$. Then, it follows from (B.6) that $\text{dist}(\tilde{x}_2^k) = \|\tilde{x}_2^k\| = \Theta(\alpha_k) = \Theta(1/k^s)$. As the order is cyclic, if we apply IG to the functions with an alternative order $f_1 = \tilde{f}_2$, $f_2 = \tilde{f}_3, \ldots, f_{m-1} = \tilde{f}_m$ and $f_m = \tilde{f}_1$ instead the resulting iterates are $x_1^k = \tilde{x}_2^k$ satisfying $\text{dist}(x_1^k) = \text{dist}(\tilde{x}_2^k) = \Theta(1/k^s)$. We conclude.

REFERENCES

[1] F. Bach. Adaptivity of averaged stochastic gradient descent to local strong convexity for logistic regression. *The Journal of Machine Learning Research*, 15(1):595–627, January 2014.

[2] D. Bertsekas. Incremental least squares methods and the extended kalman filter. *SIAM Journal on Optimization*, 6(3):807–822, 1996.

[3] D. Bertsekas. A hybrid incremental gradient method for least squares. *SIAM Journal on Optimization*, 7:913–926, 1997.

[4] D. Bertsekas. *Nonlinear programming*. Athena Scientific, 1999.

[5] D. Bertsekas. Incremental gradient, subgradient, and proximal methods for convex optimization: a survey. *Optimization for Machine Learning*, 2010:1–38, 2011.

[6] D. Bertsekas. *Convex Optimization Algorithms*. Athena Scientific, 2015.

[7] D. Bertsekas and J. Tsitsiklis. Gradient convergence in gradient methods with errors. *SIAM Journal on Optimization*, 10(3):627–642, 2000.

[8] L. Bottou and Y. Le Cun. On-line learning for very large data sets. *Applied Stochastic Models in Business and Industry*, 21(2):137–151, 2005.

[9] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends in Machine Learning*, 3(1):1–122, 2011.

[10] K. L. Chung. On a Stochastic Approximation Method. *Annals of Mathematical Statistics*, 25(3):463–483, September 1954.

[11] A. A. Gaivoronski. Convergence properties of backpropagation for neural nets via theory of stochastic gradient methods. part 1. *Optimization Methods and Software*, 4(2):117–134, 1994.

[12] L. Grippo. A class of unconstrained minimization methods for neural network training. *Optimization Methods and Software*, 4(2):135–150, 1994.

[13] M. Gürbüzbalaban, A. Ozdaglar, and P. Parrilo. A globally convergent incremental Newton method. *Mathematical Programming*, 151(1):283–313, 2015.

[14] M. Gürbüzbalaban, A. Ozdaglar, and P. Parrilo. Why random reshuffling beats stochastic gradient descent. *In preparation*, 2015.

[15] E. Hazan, A. Agarwal, and S. Kale. Logarithmic regret algorithms for online convex optimization. *Machine Learning*, 69(2-3):169–192, 2007.

[16] T. Kohonen. An adaptive associative memory principle. *IEEE Transactions on Computers*, 23(4):444–445, 1974.

[17] Z. Luo. On the convergence of the lms algorithm with adaptive learning rate for linear feedforward networks. *Neural Computation*, 3(2):226–245, 1991.

[18] Z. Luo and P. Tseng. Analysis of an approximate gradient projection method with applications to the backpropagation algorithm. *Optimization Methods and Software*, January 2008.

[19] O.L. Mangasarian and M.V. Solodov. Serial and parallel backpropagation convergence via nonmonotone perturbed minimization. *Optimization Methods and Software*, 4(2):103–116, 1994.

[20] H. Moriyama, N. Yamashita, and M. Fukushima. The incremental gauss-newton algorithm with adaptive stepsize rule. *Computational Optimization and Applications*, 26(2):107–141, 2003.

[21] E. Moulines and F. R. Bach. Non-Asymptotic Analysis of Stochastic Approximation Algorithms for Machine Learning. *Advances in Neural Information Processing*, pages 451–459, 2011.

[22] A. Nedić and D. Bertsekas. Convergence rate of incremental subgradient algorithms. In Stanislav Uryasev and PanosM. Pardalos, editors, *Stochastic Optimization: Algorithms and Applications*, volume 54 of *Applied Optimization*, pages 223–264. Springer US, 2001.

[23] A. Nedić and A. Ozdaglar. On the rate of convergence of distributed subgradient methods for multi-agent optimization. In *Proceedings of IEEE CDC*, pages 4711–4716, 2007.

[24] A. Nedic and A. Ozdaglar. Distributed subgradient methods for multi-agent optimization. *IEEE Transactions on Automatic Control*, 54(1):48–61, 2009.

[25] A. Nemirovski, A. Juditsky, G. Lan, and A. Shapiro. Robust Stochastic Approximation Approach to Stochastic Programming. *SIAM Journal on Optimization*, 19(4):1574–1609, January 2009.

[26] Y. Nesterov. *Introductory lectures on convex optimization: A basic course*, volume 87. Springer, 2004.

[27] E. Polak. On the mathematical foundations of nondifferentiable optimization in engineering design. *SIAM Review*, 29:21–90, 1987.

[28] A. Rakhlin, O. Shamir, and K. Sridharan. Making gradient descent optimal for strongly convex stochastic optimization. *arXiv preprint arXiv:1109.5647*, Sep 2011.

[29] S.S. Ram, A. Nedic, and V.V. Veeravalli. Stochastic incremental gradient descent for estimation

in sensor networks. In *Signals, Systems and Computers, ACSSC 2007.*, pages 582–586, 2007.

[30] J. Sohl-Dickstein, B. Poole, and S. Ganguli. Fast large-scale optimization by unifying stochastic gradient and quasi-Newton methods. In T. Jebara and E. P. Xing, editors, *ICML*, pages 604–612. JMLR Workshop and Conference Proceedings, 2014.

[31] M.V. Solodov. Incremental gradient algorithms with stepsizes bounded away from zero. *Computational Optimization and Applications*, 11(1):23–35, 1998.

[32] E.R. Sparks, A. Talwalkar, V. Smith, J. Kottalam, P. Xinghao, J. Gonzalez, M.J. Franklin, M.I Jordan, and T. Kraska. MLI: An API for distributed machine learning. In *IEEE 13th International Conference on Data Mining (ICDM)*, pages 1187–1192, 2013.

[33] P. Tseng. An incremental gradient(-projection) method with momentum term and adaptive stepsize rule. *SIAM Journal on Optimization*, 8(2):506–531, 1998.

[34] B. Widrow and M. E. Hoff. Adaptive switching circuits. *Institute of Radio Engineers, Western Electronics Show and Convention*, Part 4:96–104, 1960.